

11/10/99

JCS64 U.S. PTO
11/10/99

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

A

JCS64 U.S. PTO
09/437169
11/10/99

Patent Application of: Mark D. Rustad
SYSTEMS, DEVICES, STRUCTURES, AND METHODS TO SHARE RESOURCES AMONG
ENTITIES

Attorney Docket No.: 977.029US1

PATENT APPLICATION TRANSMITTAL

BOX PATENT APPLICATION

Assistant Commissioner for Patents
Washington, D.C. 20231

We are transmitting herewith the following attached items and information (as indicated with an "X"):

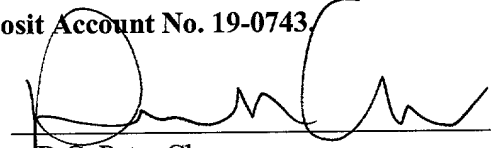
- X Utility Patent Application under 37 CFR § 1.53(b) comprising:
X Specification (28 pgs, including claims numbered 1 through 79 and a 1 page Abstract).
X Formal Drawing(s) (6 sheets).
X Signed Combined Declaration and Power of Attorney (3 pgs).
X Check in the amount of \$2,446.00 to pay the filing fee.
X Assignment of the invention to Digi International Inc. (2 pgs) and Recordation Form Cover Sheet.
X Check in the amount of \$40.00 to pay the Assignment recording fee.
X Return postcard.

The filing fee has been calculated below as follows:

	No. Filed	No. Extra	Rate	Fee
TOTAL CLAIMS	79 - 20 =	59	x 18 =	\$1,062.00
INDEPENDENT CLAIMS	11 - 3 =	8	x 78 =	\$624.00
[] MULTIPLE DEPENDENT CLAIMS PRESENTED				\$0.00
BASIC FEE				\$760.00
TOTAL				\$2,446.00

Please charge any additional required fees or credit overpayment to Deposit Account No. 19-0743.

SCHWEGMAN, LUNDBERG, WOESSNER & KLUTH, P.A.
P.O. Box 2938, Minneapolis, MN 55402 (612-373-6900)

By: 
Atty: D.C. Peter Chu
Reg. No. 41,676

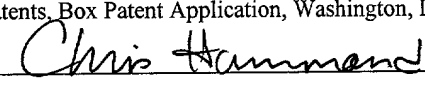
Customer Number 21186

"Express Mail" mailing label number: EL273420605US

Date of Deposit: November 3, 1999

I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to the Assistant Commissioner for Patents, Box Patent Application, Washington, D.C. 20231.

By: Chris Hammond

Signature: 

UNITED STATES PATENT APPLICATION

**SYSTEMS, DEVICES, STRUCTURES, AND METHODS TO
SHARE RESOURCES AMONG ENTITIES**

INVENTOR

Mark D. Rustad

Schwegman, Lundberg, Woessner, & Kluth, P.A.
1600 TCF Tower
121 South Eighth Street
Minneapolis, Minnesota 55402
ATTORNEY DOCKET 977.029US1
DIGI INTERNATIONAL INC.

SYSTEMS, DEVICES, STRUCTURES, AND METHODS
TO SHARE RESOURCES AMONG ENTITIES

Technical Field

5 The present invention relates generally to computer systems. More particularly, it pertains to sharing resources among a plurality of entities in computer systems.

Background Information

10 The information technology of today has grown at an unprecedented rate as a result of the synergistic marriage of communication networks and the computer. Milestones in the development of these communication networks have included the telephone networks, radio, television, cable, and communication satellites. Computers have made tremendous progress from being a single, hulking machine operated by a human operator to today's postage-stamp-size integrated circuits. The merging of the
15 communication networks and the computer has replaced the model of forcing workers to bring their work to the machine to a model of allowing anyone to access information on any computers at diverse locations and times.

 Certain barriers exist for the continuing advancement of communication networks. Communication networks have leveraged from the powerful processing capability of a
20 single computer processor. To increase processing throughput, multiple processors may be engaged in a parallel architecture. Whereas a single processor may access resources for processing in an orderly manner, each processor in a multiple-processor environment competes with the others for access to resources to complete its own processing workload. In this environment, a resource can be changed or altered by any of the
25 processors. Such changes by a processor, thus, could adversely affect the operation of other processors that are not privy to the change made by the controlling processor.

 Thus, systems, devices, structures, and methods are needed to allow resources to be shared in a multiple-processor environment.

Summary

The above-mentioned problems with sharing resources in a multiple-processor environment as well as other problems are addressed by the present invention and will be understood by reading and studying the following specification. Systems, devices, structures, and methods are described which allow resources to be shared in a multiple-processor environment.

In particular, an illustrative embodiment includes an exemplary system. This system includes a bus and a number of entities connected to the bus. At least one entity among the number of entities includes a memory. The system further includes a resource. At least a portion of the memory of one entity is selectively reset when the entity has access to the resource. For example, the memory of one entity is not reset if the entity is the same entity that previously controlled the shared resource.

Another illustrative embodiment includes an exemplary data structure in a machine-readable medium for allowing at least one resource to be shared in a multiple-processor environment, each processor in the multiple-processor environment including a fast memory. The data structure comprises a state for indicating that the resource is under control, and an identifier for identifying a past processor that had exclusive control of the resource.

A further illustrative embodiment includes an exemplary method for synchronizing access to at least one resource in a multiple-processor environment. The method comprises obtaining access to the at least one resource from a requesting processor, the requesting processor including a cache memory; excluding access to the resource except for the requesting processor; and resetting at least a portion of the cache memory of the requesting processor.

These and other embodiments, aspects, advantages, and features of the present invention will be set forth in part in the description which follows, and in part will become apparent to those skilled in the art by reference to the following description of the

invention and referenced drawings or by practice of the invention. The aspects, advantages, and features of the invention are realized and attained by means of the instrumentalities, procedures, and combinations particularly pointed out in the appended claims.

5

Brief Description of the Drawings

Figure 1 is a block diagram illustrating a system in accordance with one embodiment.

10 Figure 2 is a block diagram illustrating a system in accordance with one embodiment.

Figure 3 is a block diagram illustrating a system in accordance with one embodiment.

Figure 4 is a block diagram illustrating a system in accordance with one embodiment.

15 Figure 5 is a block diagram illustrating a data structure in accordance with one embodiment.

Figure 6 is a flow diagram illustrating a method for allowing resources to be shared in accordance with one embodiment.

20

Detailed Description

25 In the following detailed description of the invention, reference is made to the accompanying drawings that form a part hereof, and in which are shown, by way of illustration, specific embodiments in which the invention may be practiced. In the drawings, like numerals describe substantially similar components throughout the several views. These embodiments are described in sufficient detail to enable those skilled in the art to practice the invention. Other embodiments may be utilized and structural, logical, and electrical changes may be made without departing from the scope of the present invention.

Figure 1 is a block diagram illustrating a system in accordance with one embodiment. The system 100 includes a communication medium 102. In one embodiment, the communication medium 102 is a bus. In another embodiment, the communication medium 102 is a network.

5 This communication medium 102 allows data, address and controls to be communicated among manager 112, resource 116, and entities 104₀, 104₁, 104₂, ..., and 104_N. In one embodiment, at least one entity among entities 104₀, 104₁, 104₂, ..., and 104_N is an integrated circuit. In one embodiment, these entities 104₀, 104₁, 104₂, ..., and 104_N may be optionally connected to an arbiter 118 through a connection medium 114;
10 the arbiter 118 arbitrates bus requests.

Each entity 104₀, 104₁, 104₂, ..., and 104_N independently may need to access resource 116 to accomplish its workload. To access resource 116, each entity 104₀, 104₁, 104₂, ..., and 104_N obtains authorization from the manager 112. Manager 112 decides which particular entity has the authorization to access the resource 116. In one
15 embodiment, once the entity that has access to the resource 116 has accomplished its task involving the resource 116, the entity notifies the manager 112 to free up the resource 116 for other entities to use. In another embodiment, the manager 112 determines if the entity that has access to resource 116 no longer needs to use the resource 116; in this case, the manager 112 frees up the resource 116 and makes it available for other entities to access.

20 In one embodiment, the decision to grant authorization to the resource 116 is based on an algorithm. In another embodiment, the decision to grant authorization to the resource 116 is based on the priority of the workload. In another embodiment, the decision to grant authorization to the resource 116 is based upon the earliest request by an entity to access resource 116. In yet another embodiment, the decision to grant
25 authorization to the resource 116 is done in a round-robin fashion.

The manager 112, in one embodiment, is a software application such as a resource scheduler. In another embodiment, the manager 112 is an integrated circuit.

Figure 2 is a block diagram illustrating a system in accordance with one

embodiment. Integrated circuit 200₀ includes an internal bus 202. The internal bus 202 allows data and controls to be routed to and from central computing unit 204, and port controller 208.

5 The central computing unit 204 can access fast memory 222₀. In one embodiment, fast memory 222₀ is primary cache memory.

10 The port controller (or communications channel controller) 208 is receptive to communication channels 206₀, 206₁, 206₂, ... 206_N. In one embodiment at least one of these communication channels supports an asynchronous protocol. In another embodiment at least one of these communication channels supports a synchronous protocol. In another embodiment at least one of these communication channels can support either an asynchronous or a synchronous protocol. In another embodiment at least one of these communication channels supports an Asynchronous Transfer Mode (ATM) protocol. In another embodiment at least one of these communication channels supports an asymmetric digital subscriber line (ADSL) ATM protocol. In another embodiment at least one of these communication channels supports a High Level Data Link Control (HDLC) protocol. In yet another embodiment at least one of these communication channels supports transparent mode protocol. In a further embodiment, at least each of the cited protocols is controllable contemporaneously.

20 The port controller 208 manages data from communication channels 206₀, 206₁, 206₂, ... 206_N before the data is made available to the rest of integrated circuit 200₀ for further processing. The port controller 208 communicates the data from communication channels 206₀, 206₁, 206₂, ... 206_N through the internal bus 202. The port controller 208, in one embodiment, includes local area network (LAN) support. In another embodiment, the port controller 208 includes metropolitan area network (MAN) support. In yet another embodiment, the port controller 208 includes wide area network (WAN) support. In a further embodiment, the port controller 208 includes Internet support.

Interface 210 coordinates data and controls from the integrated circuit 200₀ to the bus 214. When the integrated circuit 200₀ requires access to a resource outside of the

integrated circuit 200₀, the integrated circuit 200₀ communicates with the interface 210 to establish access. When the integrated circuit 200₀ is providing data to a requesting client outside of the integrated circuit 200₀, the integrated circuit 200₀ communicates with the interface 210 to push the data to the requesting client. For illustrative purposes, the requesting client may be processor 200₁.

The bus 214 allows data and controls to be routed to and from integrated circuit 200₀, a resource 216, a processor 200₁, and a processor 200₂. In one embodiment, the resource 216 is a random access memory, such as synchronous dynamic random access memory (SDRAM). In another embodiment, the resource 216 is a memory device, such as a hard disk. In another embodiment, the resource 216 is a modifiable data source containing a data structure 218. In yet another embodiment, the resource 216 is a writable CD-ROM. In a further embodiment, the resource 216 is a computer, such as a server.

In one embodiment, the processor 200₁ includes the architecture of the integrated circuit 200₀. The processor 200₁ includes a primary fast memory 222₁. The primary fast memory 222₁ stores computer instructions and data before they are loaded into the processor 200₁ for processing. Processor 200₁ accesses the primary fast memory 222₁ for instructions and data that are needed repeatedly for program execution. The time for access to instructions and data in the primary fast memory 222₁ is shorter in comparison to the instructions and data store in secondary fast memory 224 and main memory 226₀. In one embodiment, the primary fast memory 222₁ includes cache memory.

Processor 200₂ may be similar to processor 200₁ described above. Processor 200₂ also contains primary fast memory 222₂. In one embodiment, processor 200₂ does not have secondary fast memory; instead, processor 200₂ is coupled directly to the main memory 226₁.

The switching mechanism 212 indicates whether the resource 216 is available for use. In one embodiment, the switching mechanism may be a component of the integrated circuit 200₀; in this embodiment, the switching mechanism is coupled to the central

computing unit 204. In another embodiment, the switch mechanism may be a part of the arbiter 118 of figure 1.

For illustrative purposes, suppose the processor 200₁ needs to access the resource 216 to use a portion of the resource 216. In one embodiment, the resource 216 may be a memory device. In one embodiment, the portion of resource 216 is the data structure 218. The processor 200₁ communicates with the switching mechanism 212 through the bus 214. The processor 200₁ requests the switching mechanism 212 to have access to the resource 216. If the switching mechanism 212 determines that the resource 216 is available for access, it switches control of the resource 216 to the processor 200₁. While the processor 200₁ has access to the resource 216, the switching mechanism 212 denies access to other processors that request access to the resource 216, such as processor 200₂.

At least a portion of the fast memory 222₁ may be reset when the processor 200₁ has access to the resource 216. For illustrative purposes only, suppose the processor 200₁ uses the data structure 218 repeatedly, the fast memory 222₁ stores a copy of at least a portion of the data structure 218. This will enable the processor 200₁ to spend more of its time processing since attempting to access the data structure 218 through bus 214 and resource 216 is slower than accessing a copy of the data structure through the fast memory 222₁. Once the processor 200₁ no longer needs to use the data structure 218, it informs the switching mechanism 212.

For further illustrative purposes, suppose the processor 200₂ requests access to the resource 216 to use the data structure 218. Since the resource 216 is available, the processor 200₂ obtains access to the data structure 218. A portion of the primary fast memory 222₂ of processor 200₂ is selectively reset. In one embodiment, it is understood that selectively resetting means to reset every time. Yet, in another embodiment, it is understood that selectively resetting means to reset if a certain condition is satisfied; one condition, for example, may include instances where a different processor than processor 200₂ had exclusive use of the resource 216. Processor 200₂ then makes changes to the data structure 218. Once the processor 200₂ no longer needs the data structure 218, it

informs the switching mechanism 212.

Next, for further illustrative purposes, suppose the processor 200₁ again requests access to the resource 216 to use the data structure 218. Processor 200₁ is granted access since no other processor is using the data structure 218. Processor 200₁ then proceeds to
5 access the data structure 218. But it has a copy of at least a portion of the data structure 218 in its primary fast memory 222₁ already. However, this copy of the portion of the data structure 218 may not be the same as the portion of the data structure 218 in resource 216. The portion of the data structure 218 in resource 216 may have been changed previously by processor 200₂.

10 In one embodiment, at least a portion of the fast memory 222₁ is reset so that the processor 200₁, instead of using the copy of the portion of the data structure 218 in its primary fast memory 222₁, would have to again access the portion of the data structure 218 from resource 216. In another embodiment, all of the fast memory 222₁ is reset. In one embodiment, the fast memory 222₁ is reset if another processor had access to the fast
15 memory 222₁ since the last time processor 200₁ had access to the fast memory 222₁.

In one embodiment, the switching mechanism 212 is a hardware device, such as a register. In another embodiment, the switching mechanism 212 is a software switch. In another embodiment, the switching mechanism 212 is a Dijkstra primitive. In yet another embodiment, the switching mechanism 212 may reset at least a portion of the fast
20 memory 222₁ upon granting access. In a further embodiment, the switching mechanism 212 may reset all of the fast memory 222₁ upon granting access.

Figure 3 is a block diagram illustrating a system in accordance with one embodiment. Figure 3 contains similar elements of figure 2 except that figure 3 includes an operating system 328 and a lock 330. The description of similar elements in figure 2
25 is incorporated here in figure 3. The operating system 328 is executed on the central computing unit 304. In one embodiment, the operating system 328 includes the lock 330. In another embodiment, the lock 330 exists outside the operating system 328 or outside of the integrated circuit 300₀.

0943169-4400

The lock 330 secures the resource 316 for the exclusive use of a processor. In an exemplary embodiment, the processor 300₁ needs to access the resource 316 to use a portion of the data structure 318. The processor 300₁ requests the lock 330 for exclusive access to the resource 316. If the lock 330 has not secured the resource 316 for another entity to use, it locks the resource 316 to the exclusive use of the processor 300₁. Other processors that request to use the resource 316, such as processor 300₂, wait until the resource 316 is again made available by the lock 330.

At least a portion of the fast memory 322₁ may be reset when the processor 300₁ has exclusive access to the resource 316. For illustrative purposes only, suppose the processor 300₁ uses the data structure 318 repeatedly, the fast memory 322₁ stores a copy of at least a portion of the data structure 318 to reduce bandwidth usage and memory latency. Once the processor 300₁ no longer needs to use the data structure 318, it informs the lock 330 to unlock the resource 316 for other entities to use.

For further illustrative purposes, suppose the processor 300₂ requests exclusive access to the resource 316. Since the resource 316 is available, the processor 300₂ obtains a lock to the data structure 318 for its exclusive use. A portion of the primary fast memory 322₂ of processor 300₂ is selectively reset. In one embodiment, the portion of the memory is reset every time. In another embodiment, the portion of the memory is reset if a certain condition is satisfied; one condition, for example, may include instances where a different processor than processor 300₂ had exclusive use of the resource 316. Processor 300₂ then makes changes to the data structure 318. Once the processor 300₂ no longer needs the data structure 318, it informs the lock 330 to unlock the resource 316.

Next, for further illustrative purposes, suppose the processor 300₁ again requests exclusive access to the resource 316 to use a portion of the data structure 318. Processor 300₁ obtains the lock since no other processor is using the data structure 318. Processor 300₁ then proceeds to access the portion of the data structure 318. But it has a copy of at least a portion of the data structure 318 in its primary fast memory 322₁ already. However, this copy of the portion of the data structure 318 is not the same as the portion

of the data structure 318 in resource 316. The data structure 318 in resource 316 may have been changed previously by processor 300₂. In one embodiment, at least a portion of the fast memory 322₁ is reset so that the processor 300₁, instead of using the copy of the data structure 318 in its primary fast memory 322₁, would have to again access the data structure 318 from resource 316. In another embodiment, all of the fast memory 322₁ is reset. The fast memory 322₁ is reset if another processor had access to the fast memory 322₁ since the last time processor 300₁ had access to the fast memory 322₁.

In one embodiment, the lock 330 is a hardware register. In another embodiment, the lock 330 is a software semaphore. In another embodiment, the lock 330 is a binary semaphore. In another embodiment, the lock 330 is a counting semaphore.

Figure 4 is a block diagram illustrating a system in accordance with one embodiment. Figure 4 contains similar elements of figure 3 and figure 2. The description of similar elements is incorporated here in full. Figure 4 contains together the switching mechanism 412 and the lock 430. The switching mechanism 412 may operate differently than as described heretofore.

In an illustrative embodiment, suppose the processor 400₁ needs to use a portion of the data structure 418. The processor 400₁ requests the switching mechanism 412 to grant control of the resource 416. If the switching mechanism 412 has not granted control to another processor, the switching mechanism 412 switches control to the processor 400₁. While the processor 400₁ has control, the switching mechanism 412 denies access to other processors that request similar control, such as processor 400₂. Once the processor 400₁ obtains control of the resource 416, it communicates with the switching mechanism 412 that it has control; the switching mechanism then again allow other processors to request control to the resource 416. Thus, the function of the switching mechanism 412 can be likened to a global switch.

After eliminating contention access from other processors, the processor 400₁ verifies with the lock 430 to determine whether the data structure 418 has actually been locked. In one embodiment, the lock 430 may reside within the data structure 418.

If the data structure 418 has not been locked, the processor 400₁ obtains the lock. At least a portion of the fast memory 422₁ may be reset when the processor 400₁ has access to the data structure 418. For illustrative purposes, suppose the processor 400₁ uses the data structure 418 repeatedly, the fast memory 422₁ stores a copy of at least a
5 portion of the data structure 418. From then on, the processor 400₁ uses the copy of the data structure 418 unless changes are made to the data structure 418. Once the processor 400₁ no longer needs to use the data structure 418, it informs the lock 430 and releases the lock 430 on the data structure 418.

For further illustrative purposes, suppose the processor 400₂ requests control of
10 the data structure 418 while the data structure 418 is locked by processor 400₁. Since the switching mechanism 412 has not allocated that control, the processor 400₂ obtains control. The processor 400₂ proceeds to lock the data structure 418. However, since the data structure 418 has already been locked by processor 400₁, the attempt by the processor 400₂ to lock the data structure 418 is denied. The processor 400₂ then releases
15 control to the switching mechanism 412, other processors then attempt to grab control, and the processor 400₂ waits for its chance to gain control again.

For further illustrative purposes, suppose subsequently that the processor 400₂ obtains control and gains access to the data structure 418 when the processor 400₁ releases the lock on the data structure 418. A portion of the primary fast memory 422₂ of
20 processor 400₂ is selectively reset. In one embodiment, the memory is reset every time. In another embodiment, the memory is reset when a certain condition is satisfied; one condition, for example, may include instances where a different processor than processor 400₂ had exclusive use of the resource 416. Processor 400₂ then makes changes to the data structure 418. Once the processor 400₂ no longer needs the data structure 418, it
25 releases the lock on the data structure 418.

Next, for further illustrative purposes, suppose the processor 400₁ again requests access to the data structure 418 through the above-described process. Processor 400₁ is granted access since no other processor is using the data structure 418. Processor 400₁

then proceeds to access the data structure 418. But it has a copy of at least a portion of the data structure 418 in its primary fast memory 422₁ already. However, this copy of the data structure 418 is not the same as the data structure 418 in resource 416. The data structure 418 in resource 416 may have been modified previously by processor 400₂. In one embodiment, at least a portion of the fast memory 422₁ may be reset so that the processor 400₁, instead of using the copy of the data structure 418 in its primary fast memory 422₁, would have to again access the data structure 418 from resource 416. In another embodiment, all of the fast memory 422₁ may be reset. In one embodiment, the lock 430 is a data structure. The fast memory 422₁ is only reset if another processor had access to memory since the last time processor 400₁ had access to the fast memory 422₁.

Figure 5 is a block diagram illustrating a data structure in accordance with one embodiment. Data structure 500 is used to schedule accesses to resource 516. The data structure 500 includes several data variables. The data variable "lock state" 502 contains information about whether the resource 516 has been locked or not. The data variable "last user id" 504 contains information to identify the last processor that accessed the resource 516. The data variable "present user id" 506 contains information to identify the current processor that accesses the resource 516. The data variable "resource" 524 contains at least a location and a dimension of a portion of resource 516 where such is being accessed. In one embodiment, data variable "resource" 524 is a pointer to a list 526 containing at least a location and at least a dimension of a portion of resource 516 where such is being accessed. The data variable "resource location" 508 contains the address of a portion of the resource 516. The data variable "resource dimension" 510 contains the size of a portion of the resource 516. In one embodiment, the list 526 may be implemented as an array data structure. In another embodiment, the list 526 may be implemented as a linked list.

In one embodiment, the data variables "resource location" 508 and "resource dimension" 510 are indicative of the area of the resource that the data structure 500 protects. When control of the access to the resource 516 changes hands from one

processor to another, the processor having present access resets the portion of the fast memory that relates to the area indicated by the data variables "resource location" 508 and "resource dimension" 510. In another embodiment, the data variables "resource location" 508 and "resource dimension" 510 are not used; instead, all portions of the fast memory are reset except for the portion storing stack data.

In one embodiment, the data structure 500 is a class. In that embodiment, the data structure 500 further includes a method "resetting" 512. This method is used to reset the fast memory of a processor that has access to the resource 516. This method inhibits cache incoherence so that the processor does not inadvertently use a copy of old data.

In the embodiment where the data structure 500 is a class, the data structure 500 further includes a method "comparing" 514. This method compares the data variables "last user id" 504 and "present user id" 506. If the data variables "last user id" 504 and "present user id" 506 are the same, then the method "resetting" 512 is not executed. If, however, the data variables "last user id" 504 and "present user id" 506 are different, then the method "resetting" 512 may be executed.

Processors 518_0 , 518_1 , $518_2, \dots$, and 518_N use the data structure 500 for orderly access to resource 516. The data structure 500 ensures that only one processor among processors 518_0 , 518_1 , $518_2, \dots$, and 518_N may access the resource 516 at any one time. Processors 518_0 , 518_1 , and 518_N have primary fast memory 520_0 , 520_1 , and 520_N , respectively. In one embodiment, the data structure 500 is responsible for cache coherency by resetting at least a portion of the primary fast memory of the processor currently accessing the resource 516. In another embodiment, the processors 518_0 , 518_1 , and 518_N are responsible for resetting at least a portion of the primary fast memory 520_0 , 520_1 , and 520_N , respectively, to ensure cache coherency.

In another embodiment, the processor 518_2 does not have primary fast memory. In this case, neither the data structure 500 nor the processor 518_2 needs to reset any primary fast memory.

In another embodiment, the processor 518_N has not only the primary fast memory

520_N but also secondary fast memory 522. In one embodiment, the data structure 500 would be responsible for resetting at least a portion of the primary fast memory 520_N and also at least a portion of the secondary fast memory 522 to ensure cache coherency. In another embodiment, the processor 520_N ensures cache coherency by resetting at least a portion of the primary fast memory 520_N and at least a portion of the secondary fast memory 522.

Figure 6 is a flow diagram illustrating a method for allowing resources to be shared in accordance with one embodiment. In the present embodiment, the entity that requests access to the resource to use it in some way can be a user, a processor, or a software client. For explanatory purposes, the processor will be used to describe the following embodiment.

A processor requests access to a resource to do some work. In one embodiment, such work may entail reading from the resource to obtain certain information. In another embodiment, such work may entail writing to the resource to store certain information. In another embodiment, such work may entail both reading and writing. In another embodiment, such work may be to execute certain processes on the resource. In another embodiment, such work may be to control the resource.

The processor begins by checking at block 600 to see if the resource is available. If the resource is not available, the processor then waits at block 604, and subsequently retries to gain access the resource at block 600. If the resource is available, the processor attempts to gain control of the resource at block 602. Obtaining control may include locking the resource for exclusive access and inhibiting others from contending for access.

Next, at block 604, the identity of the last processor that had accessed the resource is obtained. Then, at block 606, the identity of the present processor that has accessed the resource is obtained. At block 608, the identity of the last processor and the identity of the present processor are compared. If the identities are the same, the method goes to block 612. Otherwise, if the identities are different, block 610 resets at least a portion of

the cache memory of the processor.

At block 612, the present processor accesses the resource. Once the present processor has finished using the resource, it releases the resource so that other processors may obtain access to it.

5

Conclusion

Thus, systems, devices, structures, and methods have been described to share resources among a plurality of processors. The described embodiments allow resources to be shared without the use of complex bus snooping and cache invalidation hardware.

10 Because this hardware is also expensive, the described embodiments benefit from cost reduction. The present embodiments also enjoy an integrated solution on one chip with a small footprint because it does not use the complicated bus architecture of the bus snooping and cache invalidation hardware.

15 Although the specific embodiments have been illustrated and described herein, it will be appreciated by those of ordinary skill in the art that any arrangement which is calculated to achieve the same purpose may be substituted for the specific embodiment shown. This application is intended to cover any adaptations or variations of the present invention. It is to be understood that the above description is intended to be illustrative and not restrictive. Combinations of the above embodiments and other embodiments will
20 be apparent to those of skill in the art upon reviewing the above description. The scope of the invention includes any other applications in which the above structures and fabrication methods are used. Accordingly, the scope of the invention should only be determined with reference to the appended claims, along with the full scope of equivalents to which such claims are entitled.

25

What is claimed is:

1. A system comprising:
a bus;
5 a resource coupled to the bus; and
a plurality of entities coupled to the bus, at least one entity among the plurality of entities including a memory, wherein at least a portion of the memory of the at least one entity is selectively reset when the at least one entity has access to the resource.
- 10 2. The system of claim 1, wherein the at least one entity is an integrated circuit.
3. The system of claim 1, wherein the resource includes at least a portion of a memory device.
- 15 4. The system of claim 1, further comprising a manager to manage at least one request from the plurality of entities to access the resource.
5. The system of claim 1, further comprising an arbiter coupled to the plurality of entities to arbitrate at least one bus request from the plurality of entities.
- 20 6. The system of claim 1, wherein the at least a portion of the memory of the at least one entity is not reset when the at least one entity is the same entity that previously had control of the resource.
- 25 7. The system of claim 1, wherein the portion of the memory of the at least one entity is selectively reset when the at least one entity is different from an entity that previously had control of the resource.

8. An integrated circuit for allowing at least one resource to be controlled by a processor among a plurality of processors, at least one processor among the plurality of processors including a fast memory, the integrated circuit comprising:

a bus;

5 a central computing unit coupled to the bus; and

a switch mechanism, coupled to the central computing unit, to switch the control of the at least one resource, wherein a portion of the fast memory of at least one processor of the plurality of processors is selectively reset when the control of the at least one resource is switched.

10

9. The integrated circuit of claim 8, wherein the portion of the fast memory of the at least one processor is not reset when the at least one processor is the same processor that previously had control of the at least one resource.

15

10. The integrated circuit of claim 8, wherein the portion of the fast memory of the at least one processor is selectively reset when the at least one processor is different from a processor that previously had control of the at least one resource.

20

11. The integrated circuit of claim 8, wherein the switch mechanism is a hardware device.

12. The integrated circuit of claim 8, wherein the switch mechanism is a software switch.

25

13. The integrated circuit of claim 12, wherein the software switch is a Dijkstra primitive.

14. The integrated circuit of claim 8, wherein the at least one resource is a hardware

resource.

15. The integrated circuit of claim 14, wherein the hardware resource is a memory.

5 16. The integrated circuit of claim 8, further comprising a communications channel controller coupled to the bus.

17. The integrated circuit of claim 8, wherein the at least one resource is a software resource.

10

18. The integrated circuit of claim 17, wherein the software resource is a data structure.

19. The integrated circuit of claim 8, wherein the fast memory is cache memory.

15

20. An integrated circuit for allowing at least one resource to be shared among a plurality of processors, at least one processor of the plurality of processors including a fast memory, the integrated circuit comprising:

a bus;

20

a central computing unit coupled to the bus; and

a lock coupled to the central computing unit to reserve exclusive control of the at least one resource to a processor among the plurality of processors, wherein at least a portion of the fast memory of the processor of the plurality of processors is selectively reset when the processor of the plurality of processors has exclusive control of the at least one resource.

25

21. The integrated circuit of claim 20, wherein the portion of the fast memory of the processor is not reset when the processor is the same processor that previously had

exclusive control of the at least one resource.

22. The integrated circuit of claim 20, wherein the portion of the fast memory of the processor is selectively reset when the processor is different from another processor that previously had control of the at least one resource.

23. The integrated circuit of claim 20, wherein the lock is a hardware register.

24. The integrated circuit of claim 20, wherein the lock is a software semaphore.

25. The integrated circuit of claim 24, wherein the software semaphore is a binary semaphore.

26. The integrated circuit of claim 20, further comprising a communications channel controller coupled to the bus.

27. The integrated circuit of claim 20, wherein the fast memory is cache memory.

28. The integrated circuit of claim 27, wherein the cache memory is primary cache memory.

29. The integrated circuit of claim 27, wherein the cache memory is secondary cache memory.

30. A data structure in a machine-readable medium for allowing at least one resource to be shared among a plurality of processors, at least one processor of the plurality of processors including a fast memory, the data structure comprising:

a state for indicating that the at least one resource is under control; and

a first identifier for identifying a past processor that had exclusive control of the at least one resource.

31. The data structure of claim 30, wherein the data structure is a class, the data
5 structure further comprising an act for resetting at least a portion of the fast memory of the present processor.

32. The data structure of claim 31, further comprising a second identifier for
10 identifying the present processor that has exclusive control of the at least one resource.

33. The data structure of claim 32, further comprising an act for comparing the first
15 identifier and the second identifier, and wherein the act for resetting the fast memory of the present processor being executed when the first identifier is different from the second identifier.

34. The data structure of claim 30, wherein the fast memory is cache memory.

35. The data structure of claim 30, further comprising a data type that is adapted to
20 represent at least one portion of the at least one resource, wherein the data type includes at least one location of the at least one portion of the at least one resource and at least one dimension of the at least one portion of the at least one resource.

36. The data structure of claim 30, further comprising a list that includes at least one
25 location of at least one portion of the at least one resource and at least one dimension of at least one portion of the at least one resource.

37. A method for allowing at least one resource to be shared among a plurality of
processors, the method comprising:

obtaining exclusive control over the at least one resource by a present processor,
the present processor including a fast memory;

identifying a past processor to obtain a first identity, wherein the past processor
had exclusive control over the at least one resource; and

5 resetting selectively at least a portion of the fast memory of the present processor
when the past processor is different from the present processor.

38. The method of claim 37, wherein identifying the present processor further
comprises the fast memory as cache memory.

10

39. The method of claim 38, further comprising identifying a present processor to
obtain a second identity, the present processor having exclusive control over the at least
one resource, the present processor including a fast memory.

15 40. The method of claim 39, further comprising comparing the first identity and the
second identity so as to determine if the present processor is different from the past
processor.

20 41. The method of claim 37, wherein the progression of the method is in the order
presented.

42. A method for scheduling access to at least one resource from among a plurality of
processors, the method comprising:

25 obtaining access to the at least one resource from a requesting processor, the
requesting processor including a cache memory;

 excluding access to the at least one resource from the plurality of processors
except for the requesting processor; and

 resetting at least a portion of the cache memory of the requesting processor when

the requesting processor is different from a processor that previously had access to the at least one resource.

43. An integrated circuit for allowing at least one resource to be controlled by a processor among a plurality of processors, at least one processor among the plurality of processors including a fast memory, the integrated circuit comprising:

a bus;

a central computing unit coupled to the bus;

a switch mechanism for switching the control of the at least one resource; and

a lock, in a cooperative relationship with the switching mechanism, for reserving exclusive control of the at least one resource to a processor among the plurality of processors, wherein at least a portion of the fast memory of the processor of the plurality of processors is selectively reset when the processor of the plurality of processors has exclusive control of the at least one resource.

44. The integrated circuit of claim 43, wherein the fast memory is cache memory.

45. The integrated circuit of claim 43, further comprising a communications channel controller coupled to the bus, wherein the communications channel controller is receptive to diverse communications protocols.

46. The integrated circuit of claim 43, wherein the cooperative relationship of the switch mechanism and the lock maintains cache coherency.

47. The integrated circuit of claim 43, wherein the at least a portion of the fast memory of the at least one processor is not reset when the processor is the same processor that previously had control of the at least one resource.

48. The integrated circuit of claim 43, wherein the portion of the fast memory of the processor is selectively reset when the processor is different from a processor that previously had control of the at least one resource.

5 49. An integrated circuit for allowing at least one resource to be controlled by a processor among a plurality of processors, at least one processor among the plurality of processors including a fast memory, the integrated circuit comprising:

a bus;

a central computing unit coupled to the bus; and

10 a scheduler, coupled to the central computing unit, for scheduling the control of the at least one resource, wherein a portion of the fast memory of at least one processor of the plurality of processors is selectively reset when the at least one resource is under control.

15 50. The integrated circuit of claim 49, wherein the portion of the fast memory of the at least one processor is not reset when the at least one processor is the same processor that previously had control of the at least one resource.

20 51. The integrated circuit of claim 49, wherein the portion of the fast memory of the at least one processor is selectively reset when the at least one processor is different from a processor that previously had control of the at least one resource.

52. A system comprising:

a bus;

25 at least one resource coupled to the bus;

a plurality of processors coupled to the bus, at least one processor among the plurality of processors including a fast memory; and

a switch mechanism, coupled to the bus, to switch the control of the at least one

resource, wherein a portion of the fast memory of at least one processor of the plurality of processors is selectively reset when the control of the at least one resource is switched.

53. The system of claim 52, wherein the portion of the fast memory of the at least one
5 processor is not reset when the at least one processor is the same processor that
previously had control of the at least one resource.

54. The system of claim 52, wherein the portion of the fast memory of the at least one
processor is selectively reset when the at least one processor is different from a processor
10 that previously had control of the at least one resource.

55. The system of claim 52, wherein the switch mechanism is a hardware device.

56. The system of claim 52, wherein the switch mechanism is a software switch.
15

57. The system of claim 56, wherein the software switch is a Dijkstra primitive.

58. The system of claim 52, wherein the at least one resource is a hardware resource.

59. The system of claim 58, wherein the hardware resource is a memory.
20

60. The system of claim 52, wherein the at least one processor includes a
communications channel controller.

61. The system of claim 52, wherein the at least one resource is a software resource.
25

62. The system of claim 61, wherein the software resource is a data structure.

63. The system of claim 52, wherein the fast memory is cache memory.

64. A system comprising:

a bus;

5 a at least one resource coupled to the bus;

a plurality of processors coupled to the bus, at least one processor of the plurality of processors including a fast memory; and

a lock to reserve exclusive control of the at least one resource to a processor among the plurality of processors, wherein at least a portion of the fast memory of the processor of the plurality of processors is selectively reset when the processor of the plurality of processors has exclusive control of the at least one resource.

65. The system of claim 64, wherein the portion of the fast memory of the processor is not reset when the processor is the same processor that previously had exclusive control of the at least one resource.

66. The system of claim 64, wherein the portion of the fast memory of the processor is selectively reset when the processor is different from another processor that previously had control of the at least one resource.

67. The system of claim 64, wherein the lock is a hardware register.

68. The system of claim 64, wherein the lock is a software semaphore.

69. The system of claim 68, wherein the software semaphore is a binary semaphore.

70. The system of claim 64, further comprising a communications channel controller coupled to the bus.

- 60077-697260
71. The system of claim 64, wherein the fast memory is cache memory.
72. The system of claim 71, wherein the cache memory is primary cache memory.
- 5 73. The system of claim 71, wherein the cache memory is secondary cache memory.
74. A system comprising:
- a bus;
- at least one resource coupled to the bus;
- 10 a plurality of processors coupled to the bus, at least one processor among the plurality of processors including a fast memory;
- a switch mechanism to switch the control of the at least one resource; and
- a lock, in a cooperative relationship with the switching mechanism, to reserve exclusive control of the at least one resource to a processor among the plurality of
- 15 processors, wherein at least a portion of the fast memory of the processor of the plurality of processors is selectively reset when the processor of the plurality of processors has exclusive control of the at least one resource.
75. The system of claim 74, wherein the fast memory is cache memory.
- 20 76. The system of claim 74, wherein the at least one processor includes a communications channel controller, wherein the communications channel controller is receptive to diverse communications protocols.
- 25 77. The system of claim 74, wherein the cooperative relationship of the switch mechanism and the lock maintains cache coherency.
78. The system of claim 74, wherein the at least a portion of the fast memory of the at

SYSTEMS, DEVICES, STRUCTURES, AND METHODS TO SHARE RESOURCES
AMONG ENTITIES

5

Abstract of the Disclosure

Systems, devices, structures, and methods are provided to allow resources to be shared among a plurality of processors. An exemplary system includes a mechanism to grant exclusive control of a resource to a processor, while at the same time, the fast memory of such a processor is maintained in a coherent state. An exemplary structure
10 includes data structures that help to identify the portion of the fast memory of the processor to be maintained in a coherent state. An exemplary method includes a determination of past and present processors that have had access to the resource so as to maintain the coherency of the fast memory of the processor.

00437459-11000

"Express Mail" mailing label number: EL 27342060545
Date of Deposit: Nov. , 1999
I hereby certify that this paper or fee is being deposited with the
United States Postal Service * Express Mail Post Office to Addressee*
service under 37 CFR 1.10 on the date indicated above and is
addressed to the Assistant Commissioner for Patents,
Washington, D.C. 20231
Printed Name Chris Hammond
Signature Chris Hammond

FIG. 1




FIG. 2

66077-092460




FIG. 3





FIG. 4

6607494






FIG. 5

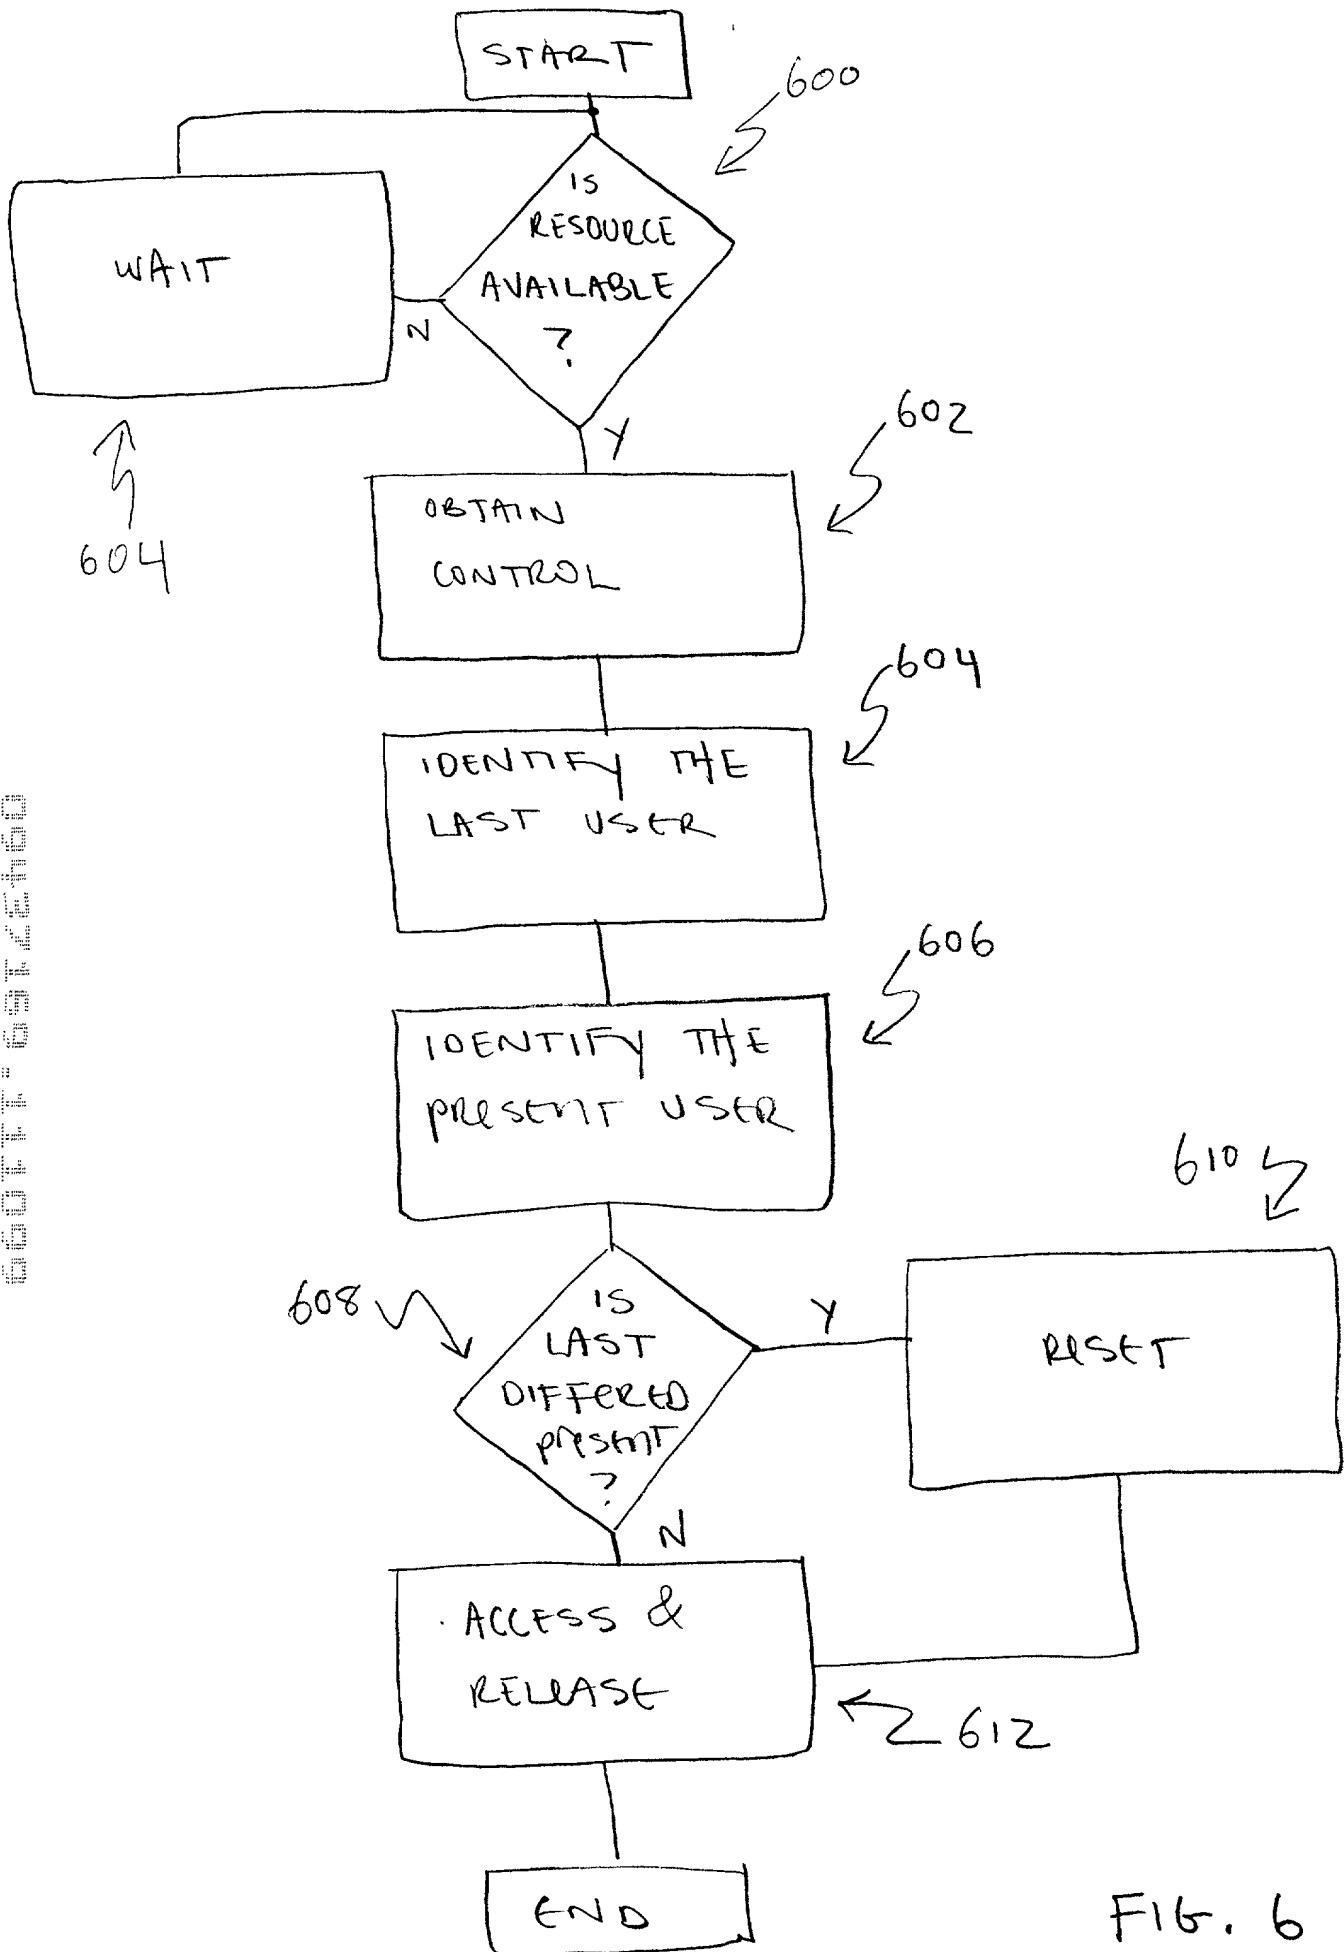


FIG. 6

SCHWEGMAN, LUNDBERG, WOESSNER & KLUTH, P.A.

United States Patent Application

COMBINED DECLARATION AND POWER OF ATTORNEY

As a below named inventor I hereby declare that: my residence, post office address and citizenship are as stated below next to my name; that

I verily believe I am the original, first and sole inventor of the subject matter which is claimed and for which a patent is sought on the invention entitled: **SYSTEMS, DEVICES, STRUCTURES, AND METHODS TO SHARE RESOURCES AMONG ENTITIES.**

The specification of which is attached hereto.

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the patentability of this application in accordance with 37 C.F.R. § 1.56 (attached hereto). I also acknowledge my duty to disclose all information known to be material to patentability which became available between a filing date of a prior application and the national or PCT international filing date in the event this is a Continuation-In-Part application in accordance with 37 C.F.R. §1.63(e).

I hereby claim foreign priority benefits under 35 U.S.C. §119(a)-(d) or 365(b) of any foreign application(s) for patent or inventor's certificate, or 365(a) of any PCT international application which designated at least one country other than the United States of America, listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on the basis of which priority is claimed:

No such claim for priority is being made at this time.

I hereby claim the benefit under 35 U.S.C. § 119(e) of any United States provisional application(s) listed below:

No such claim for priority is being made at this time.

I hereby claim the benefit under 35 U.S.C. § 120 or 365(c) of any United States and PCT international application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States or PCT international application in the manner provided by the first paragraph of 35 U.S.C. § 112, I acknowledge the duty to disclose material information as defined in 37 C.F.R. § 1.56(a) which became available between the filing date of the prior application and the national or PCT international filing date of this application:

No such claim for priority is being made at this time.

I hereby appoint the following attorney(s) and/or patent agent(s) to prosecute this application and to transact all business in the Patent and Trademark Office connected herewith:

Adams, Gregory J.	Reg. No. P-44,494	Forrest, Bradley A.	Reg. No. 30,837	McCrackin, Ann M.	Reg. No. 42,858
Adams, Matthew W.	Reg. No. 43,459	Harris, Robert J.	Reg. No. 37,346	Nama, Kash	Reg. No. 44,255
Anglin, J. Michael	Reg. No. 24,916	Huebsch, Joseph C.	Reg. No. 42,673	Nelson, Albin J.	Reg. No. 28,650
Arora, Suneel	Reg. No. 42,267	Jurkovich, Patti J.	Reg. No. P-44,813	Nielsen, Walter W.	Reg. No. 25,539
Bianchi, Timothy E.	Reg. No. 39,610	Kalis, Janal M.	Reg. No. 37,650	Oh, Allen J.	Reg. No. 42,047
Billion, Richard E.	Reg. No. 32,836	Kaufmann, John D.	Reg. No. 24,017	Padys, Danny J.	Reg. No. 35,635
Black, David W.	Reg. No. 42,331	Klima-Silberg, Catherine I.	Reg. No. 40,052	Parker, J. Kevin	Reg. No. 33,024
Brennan, Thomas F.	Reg. No. 35,075	Kluth, Daniel J.	Reg. No. 32,146	Peacock, Gregg A.	Reg. No. P-45,001
Brooks, Edward J., III	Reg. No. 40,925	Lacy, Rodney L.	Reg. No. 41,136	Polglaze, Daniel J.	Reg. No. 39,801
Chu, Dinh C.P.	Reg. No. 41,676	Leffert, Thomas W.	Reg. No. 40,697	Prout, William F.	Reg. No. 33,995
Clark, Barbara J.	Reg. No. 38,107	Lemaire, Charles A.	Reg. No. 36,198	Schwegman, Micheal L.	Reg. No. 25,816
Dahl, John M.	Reg. No. P-44,639	Litman, Mark A.	Reg. No. 26,390	Sieffert, Kent J.	Reg. No. 41,312
Drake, Eduardo E.	Reg. No. 40,594	Lundberg, Steven W.	Reg. No. 30,568	Slifer, Russell D.	Reg. No. 39,838
Eliseeva, Maria M.	Reg. No. 43,328	Mack, Lisa K.	Reg. No. 42,825	Steffey, Charles E.	Reg. No. 25,179
Embretson, Janet E.	Reg. No. 39,665	Maki, Peter C.	Reg. No. 42,832	Terry, Kathleen R.	Reg. No. 31,884
Fogg, David N.	Reg. No. 35,138	Malen, Peter L.	Reg. No. P-44,894	Viksnins, Ann S.	Reg. No. 37,748
Fordenbacher, Paul J.	Reg. No. 42,546	Mates, Robert E.	Reg. No. 35,271	Woessner, Warren D.	Reg. No. 30,440


I hereby authorize them to act and rely on instructions from and communicate directly with the person/assignee/attorney/firm/organization/who/which first sends/sent this case to them and by whom/which I hereby declare that I have consented after full disclosure to be represented unless/until I instruct Schwegman, Lundberg, Woessner & Kluth, P.A. to the contrary.

Please direct all correspondence in this case to **Schwegman, Lundberg, Woessner & Kluth, P.A.** at the address indicated below:

P.O. Box 2938, Minneapolis, MN 55402

Telephone No. (612)373-6900

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full Name of sole inventor : **Mark D. Rustad**
Citizenship: **United States of America** Residence: **Edina, MN**
Post Office Address: **6505 Polar Circle**
Edina, MN 55436
Signature:  Date: 10/8/99
Mark D. Rustad

§ 1.56 Duty to disclose information material to patentability.

(a) A patent by its very nature is affected with a public interest. The public interest is best served, and the most effective patent examination occurs when, at the time an application is being examined, the Office is aware of and evaluates the teachings of all information material to patentability. Each individual associated with the filing and prosecution of a patent application has a duty of candor and good faith in dealing with the Office, which includes a duty to disclose to the Office all information known to that individual to be material to patentability as defined in this section. The duty to disclose information exists with respect to each pending claim until the claim is canceled or withdrawn from consideration, or the application becomes abandoned. Information material to the patentability of a claim that is canceled or withdrawn from consideration need not be submitted if the information is not material to the patentability of any claim remaining under consideration in the application. There is no duty to submit information which is not material to the patentability of any existing claim. The duty to disclose all information known to be material to patentability is deemed to be satisfied if all information known to be material to patentability of any claim issued in a patent was cited by the Office or submitted to the Office in the manner prescribed by §§ 1.97(b)-(d) and 1.98. However, no patent will be granted on an application in connection with which fraud on the Office was practiced or attempted or the duty of disclosure was violated through bad faith or intentional misconduct. The Office encourages applicants to carefully examine:

- (1) prior art cited in search reports of a foreign patent office in a counterpart application, and
- (2) the closest information over which individuals associated with the filing or prosecution of a patent application believe any pending claim patentably defines, to make sure that any material information contained therein is disclosed to the Office.

(b) Under this section, information is material to patentability when it is not cumulative to information already of record or being made of record in the application, and

- (1) It establishes, by itself or in combination with other information, a prima facie case of unpatentability of a claim; or
- (2) It refutes, or is inconsistent with, a position the applicant takes in:
 - (i) Opposing an argument of unpatentability relied on by the Office, or
 - (ii) Asserting an argument of patentability.

A prima facie case of unpatentability is established when the information compels a conclusion that a claim is unpatentable under the preponderance of evidence, burden-of-proof standard, giving each term in the claim its broadest reasonable construction consistent with the specification, and before any consideration is given to evidence which may be submitted in an attempt to establish a contrary conclusion of patentability.

(c) Individuals associated with the filing or prosecution of a patent application within the meaning of this section are:

- (1) Each inventor named in the application;
- (2) Each attorney or agent who prepares or prosecutes the application; and
- (3) Every other person who is substantively involved in the preparation or prosecution of the application and who is associated with the inventor, with the assignee or with anyone to whom there is an obligation to assign the application.

(d) Individuals other than the attorney, agent or inventor may comply with this section by disclosing information to the attorney, agent, or inventor.